# IDR WG 93

# Monday Session

http://trac.tools.ietf.org/wg/idr/trac/wiki

- Status of Drafts:  14 IESG, 6 new, 5 old, 3 pending
  - 3 at IESG,  7 going to IESG, 2 Early adoption IESG
  - 5 new drafts, 3 in adoption, 4 await implementations
  - 3 Administrative drafts (2 ready for IESG)
  - Not passed WG adoption call:  3
  - Not passed WG LC: 1
- Draft authors will be responsible for
  - Protocol Implementation reports on Wiki
  - Testimonials for Administrative Drafts

# Agenda (1)

7/20/2015 17:40 - 18:40 Prague Time

Admin Trivia
==========
Agenda Bashing and Status Q&A 3 minutes
 Due to the short IDR meetings.
 The status is online.
 http://trac.tools.ietf.org/wg/idr/trac/wiki/idr-draft-status

The chairs will answer questions on status on Monday.

# Agenda (2)

**Existing Work: [30 minutes ]**

| | |
|---|---|
| draft-ietf-idr-bgpls-segment-routing-epe. | 3 minutes |
| (Stefano Previdi) | [17:40-17:45] |
| draft-ietf-idr-te-lsp-distribution | 7 minutes |
| (Jie Dong) | [17:45-17:52] |
| draft-ietf-idr-bgp-optimal-route-reflection | 5 minutes |
| (Bruno Decraene) | [17:55-18:00] |
| ietf-rs-bfd | |
| (Randy Bus) | [18:00-18:10] |
| draft-jdurand-auto-bfd-00 | 3 minutes |

# Monday Agenda (3)

Update on proposed drafts [18:20-18:50]
draft-keyupate-idr-bgp-prefix-sid 5 minutes
(Stefano Previdi) [18:10-18:15]

draft-walton-bgp-hostname-capability-00 10 minutes
(Daniel Walton) [18:15-18:25]

draft-fang-idr-bgplu-for-hsdn 15 minutes
[Luyuan Fang] [18:25-18:40]

# Friday Agenda 1

| | | |
|---|---|---|
| Agenda Bashing and status | [5 minutes ] | [Sue Haers] |
| draft-ietf-idr-flowspec-l2vpn-01 | [5 minutes ] | [Weiguo Hao] |
| draft-hao-ls-trill-01 | [10 minutes] | [Donald Eastlake] |
| | | |
| draft-hao-idr-flowspec-nv03-00 | [5 minutes] | [Weiguo Hao] |
| draft-li-idr-flowspec-rpd | [15 minutes] | [Eric Wu] |
| draft-liang-idr-bgp-flowspec-label | [10 minutes] | [Jianjie You] |
| draft-wu-idr-flowspec-yang-cfg | [10 minutes] | [Eric Wu] |
| draft-liang-idr-flowspec-orf-00 | [10 minutes] | [Weiguo Hao] |

# Dissemination of Flow Specification Rules for L2 VPN

## draft-ietf-idr-flowspec-l2vpn-01

Weiguo Hao Huawei
Qiandeng Liang Huawei
Jim Uttaro AT&T
S. Litkowski Orange
Shunwan Zhuang Huawei

July, 2015 Prague

# History

## draft-hao-idr-flowspec-l2vpn-00

Presented in 91st IETF and received some comments:

- A new SAFI for L2VPN flow-spec application is not suggested.
- Should consider more scenarios like QinQ,SNAP,etc

## draft-hao-idr-flowspec-l2vpn-01

Make revisions for the comments

- SAFI 134 is redefined for of VPN flow specification rules.
- Add additional component types for QinQ, SNAP,etc

Close to 92nd IETF: Adopted as WG draft of 'draft-ietf-idr-flowspec-l2vpn-00'

# Updates since draft-ietf-idr-flowspec-l2vpn-00

Traffic Actions:

Add VLAN action and TPID action.

| type | extended community | encoding |
|---|---|---|
| 0x800A | VLAN-action | bitmask |
| 0x800B | TPID-action | bitmask |

# Updates con. (VLAN Action)

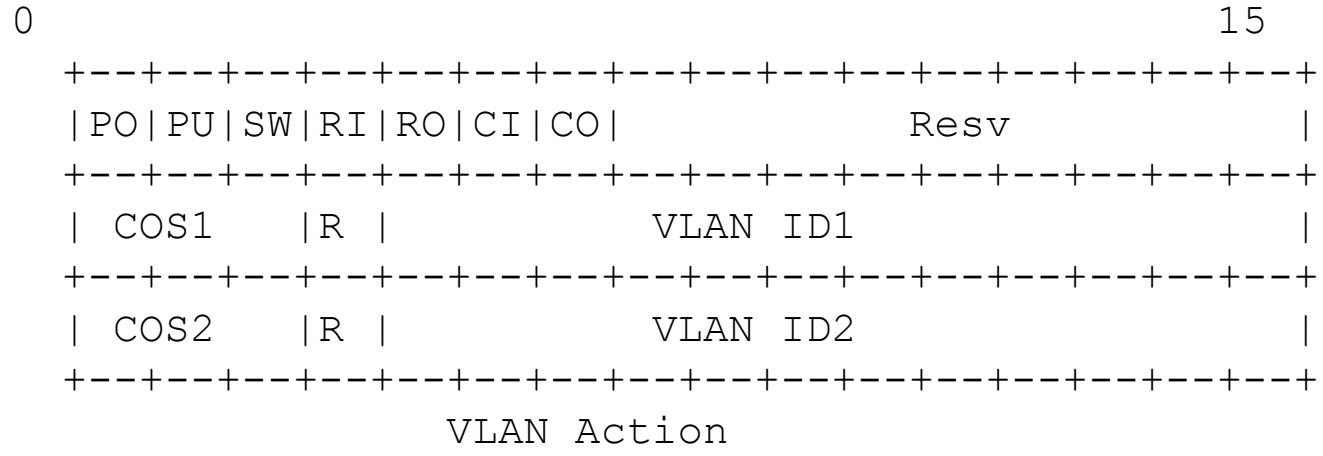


**PO:** Pop action.

**PU:** Push action.: push VLAN is VLAN ID1.

**SW:** Swap action.

**RI:** Rewrite inner VLAN action. The new VLAN is VLAN ID1.

**RO:** Rewrite outer VLAN action. The new VLAN is VLAN ID2.

**CI**: Mapping inner COS action. The new COS is COS1.

**CO:** Mapping outer COS action. The new COS is COS2.

**Resv:** Reservead for future use.

**COS1**: 3 bits. COS value.

**COS2:** 3 bits. COS value.

**VLAN ID1:** 12 bits. VLAN ID value.

**VLAN ID2:** 12 bits. VLAN ID value.

# Updates con. (TP ID Action)

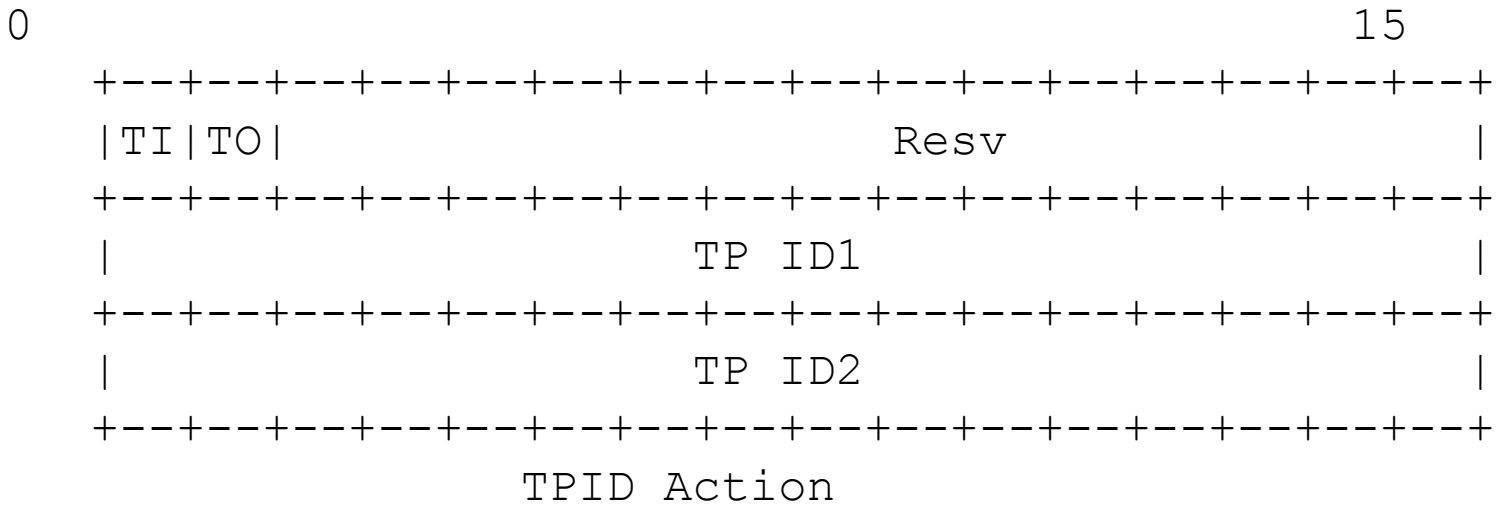


TI: Mapping inner TP ID action. The new TP ID is TP ID1.

TO: Mapping outer TP ID action. The new TP ID is TP ID2.

# Next Step

- Redirect to MAC VRF
- Flowspec to a specific interface(s)
- Other suggestions?

# Distribution of TRILL Link-State using BGP

draft-hao-idr-ls-trill-01

Weiguo Hao
Donald Eastlake
Huawei

July 2015 Prague

# Motivation

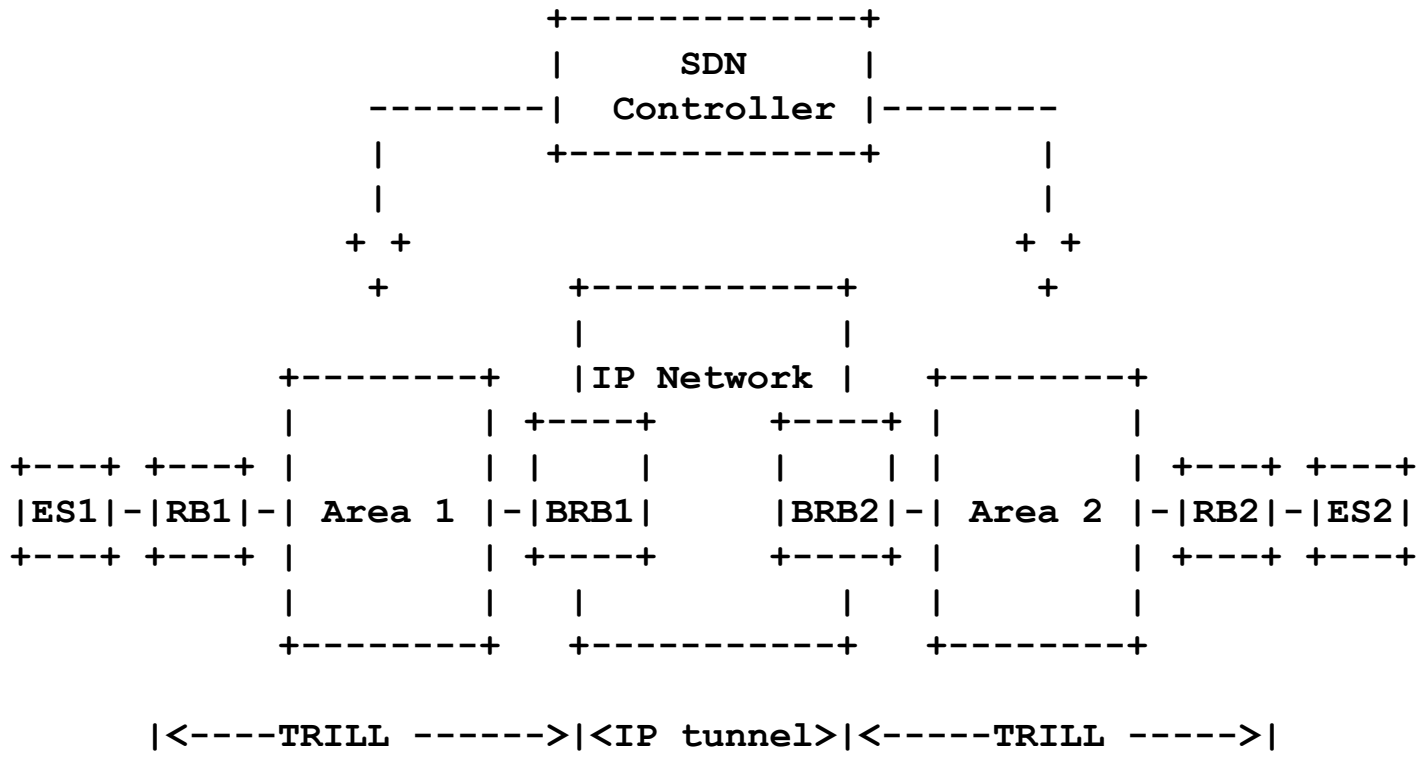


1. End-to-end topology visibility on the SDN controller
2. MAC address reachability information synchronization across multiple TRILL domains

Extension of BGP Link state is proposed to support TRILL link-state and MAC address reachability information distribution.

# New Protocol-ID for TRILL

To differentiate the TRILL protocol from layer 3 IGP protocol, a new TRILL Protocol-ID is defined.

| Protocol-ID | NLRI information source protocol |
|---|---|
| 1 | IS-IS Level 1 |
| 2 | IS-IS Level 2 |
| 3 | OSPFv2 |
| 4 | Direct |
| 5 | Static configuration |
| 6 | OSPFv3 |
| TBD | TRILL |

# New NLRI Type: MAC Reachability NLRI

| Type | NLRI Type |
|---|---|
| 1 | Node NLRI |
| 2 | Link NLRI |
| 3 | IPv4 Topology Prefix NLRI |
| 4 | IPv6 Topology Prefix NLRI |
| TBD | MAC Reachability NLRI |

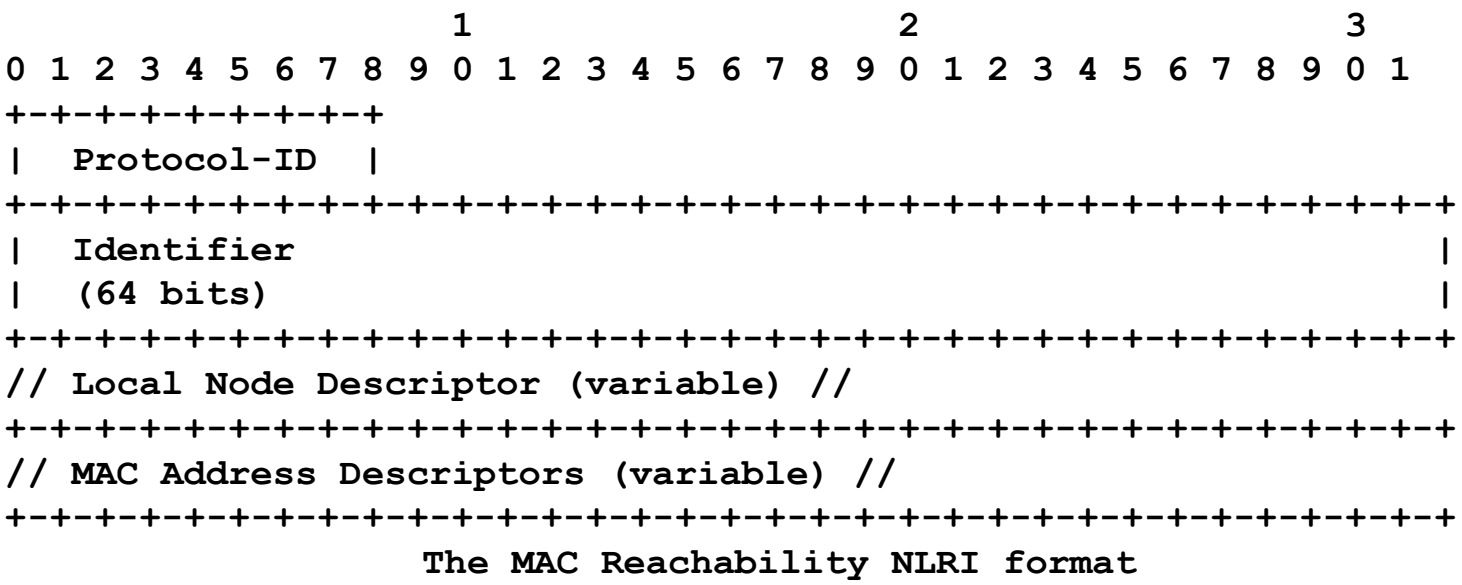


The MAC Reachability NLRI format

# MAC Address Descriptors

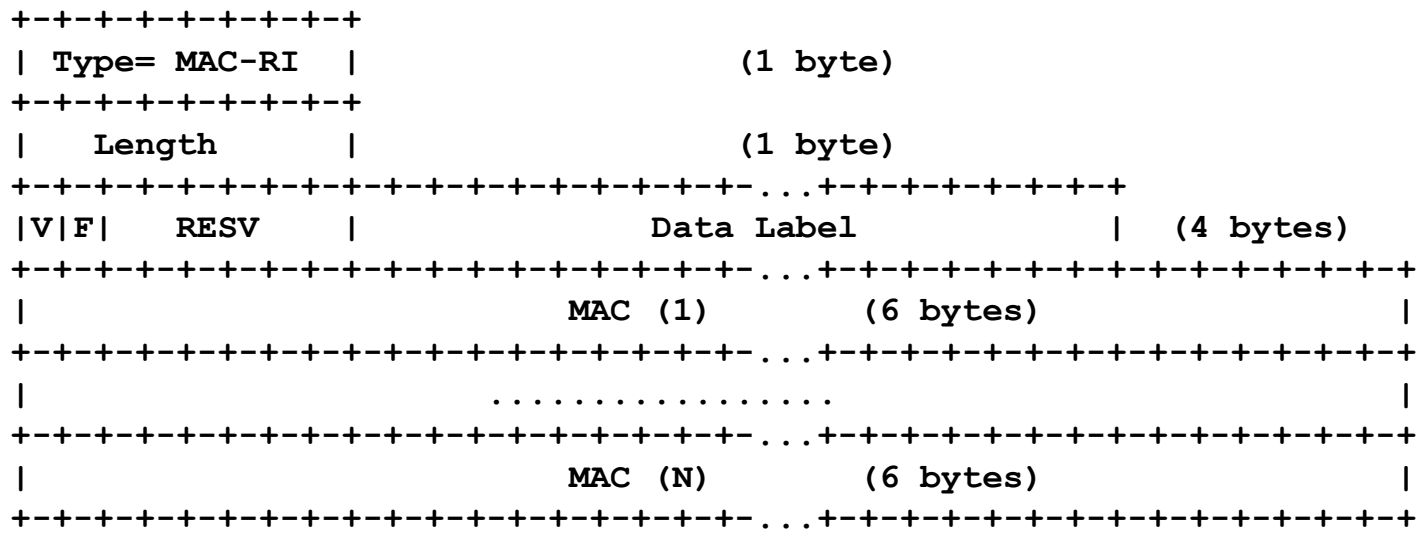


MAC-Reachability TLV format

V: VLAN

F: Fine Grained Label

# Opaque Node Attribute TLV

The Opaque Node Attribute TLV could be used as the envelope to transparently carry TRILL specific information.

| Descriptions | IS-IS TLV/Sub-TLV |
|---|---|
| TRILL-VER | 22/13 |
| NICKNAME | 22/6 |
| TREES | 22/7 |
| TREE-RT-IDs | 22/8 |
| TREE-USE-IDs | 22/9 |
| INT-VLAN | 22/10 |
| VLAN-GROUP | 22/14 |
| INT-LABEL | 22/15 |
| RBCHANNELS | 22/16 |
| AFFINITY | 22/17 |
| LABEL-GROUP | 22/18 |
| GMAC-ADDR | 142/1 |
| GIP-ADDR | 142/2 |
| GIPV6-ADDR | 142/3 |
| GLMAC-ADDR | 142/4 |
| GLIP-ADDR | 142/5 |
| GLIPV6-ADDR | 142/6 |

TRILL TLVs/Sub-TLVs

Router Capability
MT-Capability
Group Address (GADDR) TLV

# Link Attribute TLVs

| TLV Code Point | Description | IS-IS TLV /Sub-TLV | Defined in: |
|---|---|---|---|
| TBD | Link MTU | 22/28 | [RFC7176]/2.4 |

**Link Attribute TLVs**

# Node Attribute TLVs: Node Flag Bits TLV

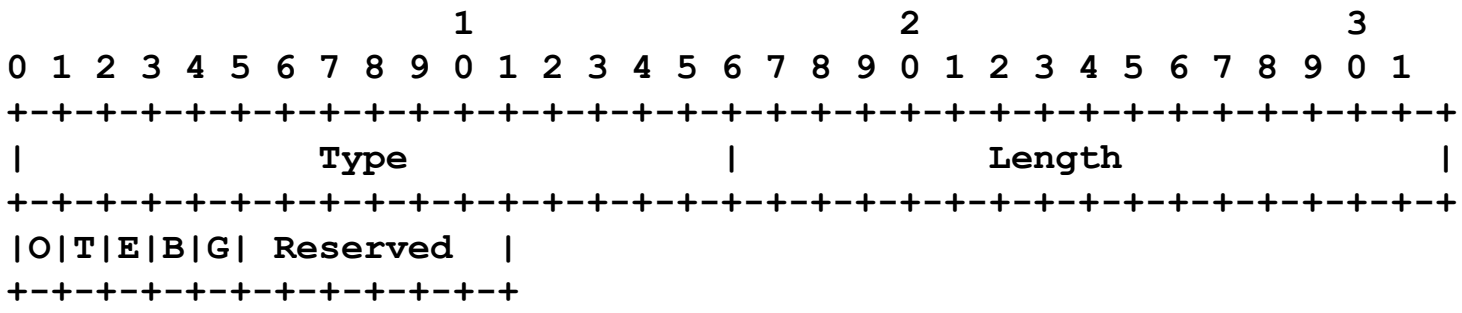


**Node Flag Bits TLV format**

**The bits are defined as follows:**

| Bit | Description | Reference |
|---|---|---|
| 'G' | Layer 3 Gateway Bit | [RFC7176] |
| Reserved | Reserved for future use | |

**Node Flag Bits Definitions**

# Next Step

- Seek some comments and feedbacks
- WG adoption?

# Dissemination of Flow Specification Rules for NVO3

draft-hao-idr-flowspec-nvo3-00

Weiguo Hao Huawei
Shunwan Zhuang Huawei
Zhenbin Li Huawei

July, 2015 Prague

# Scenario 1: DDOS mitigation

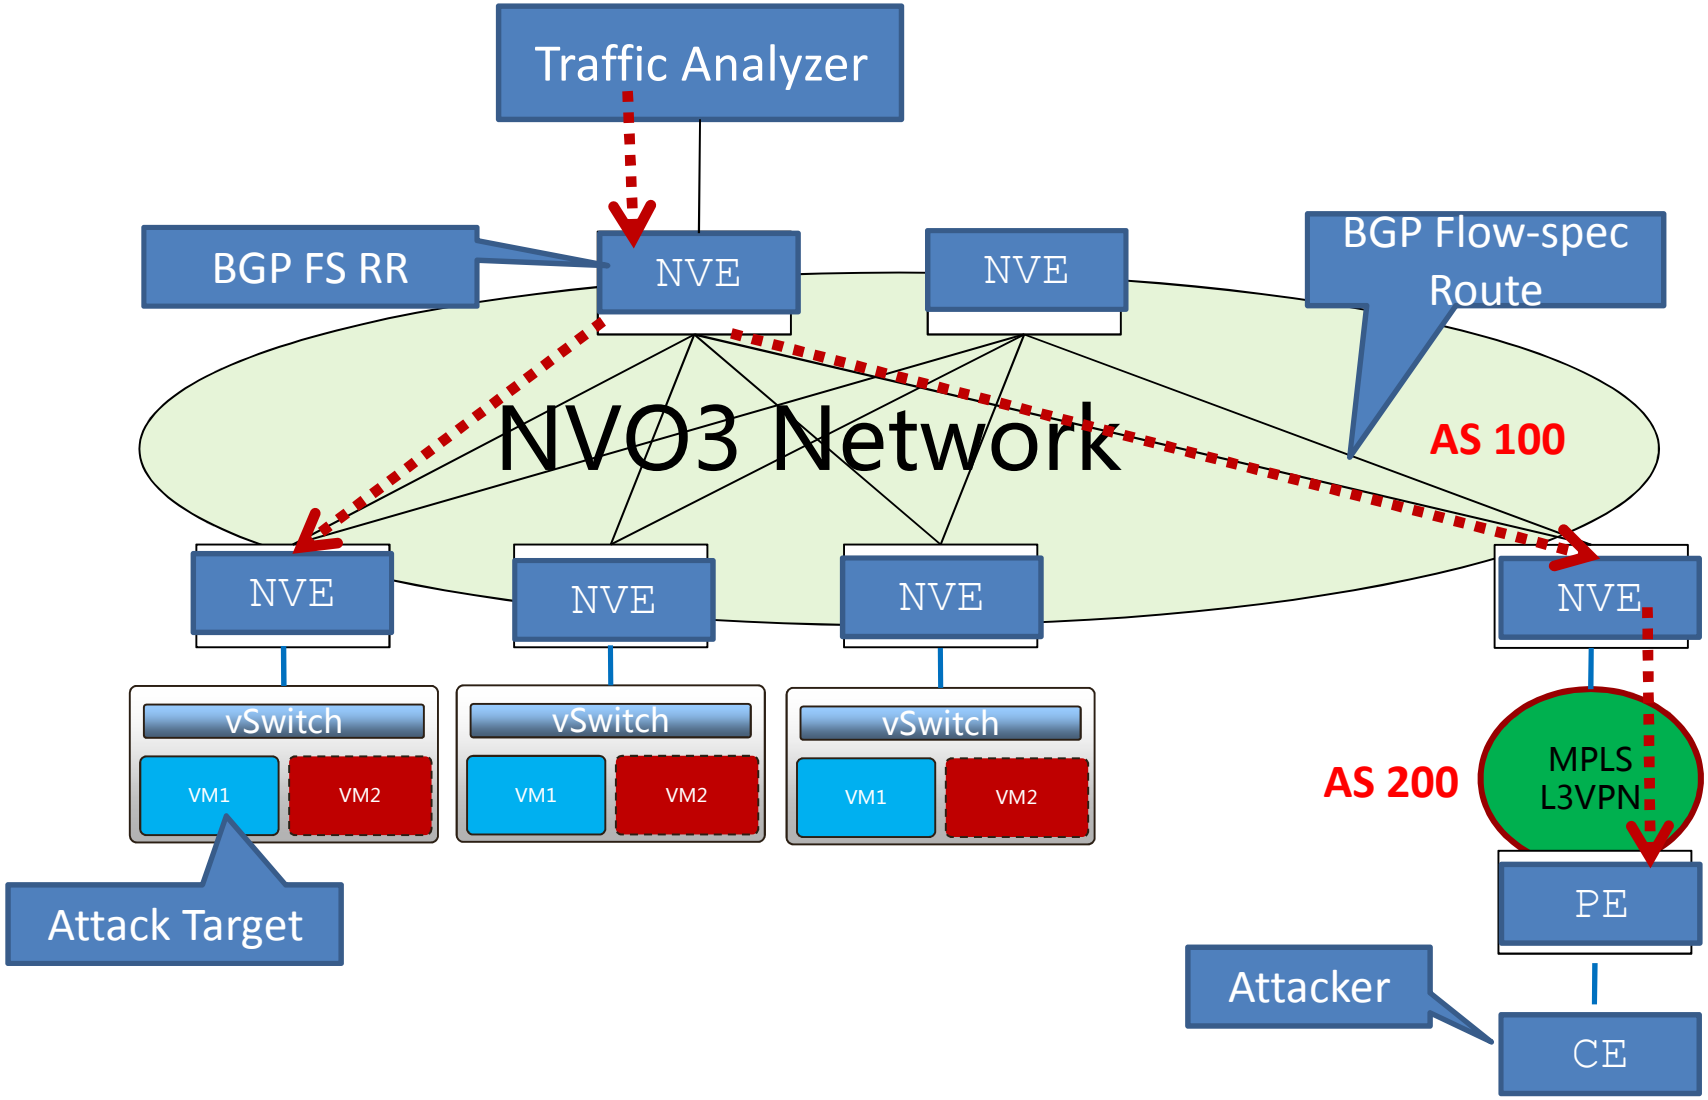

# Scenario 2: Traffic steering

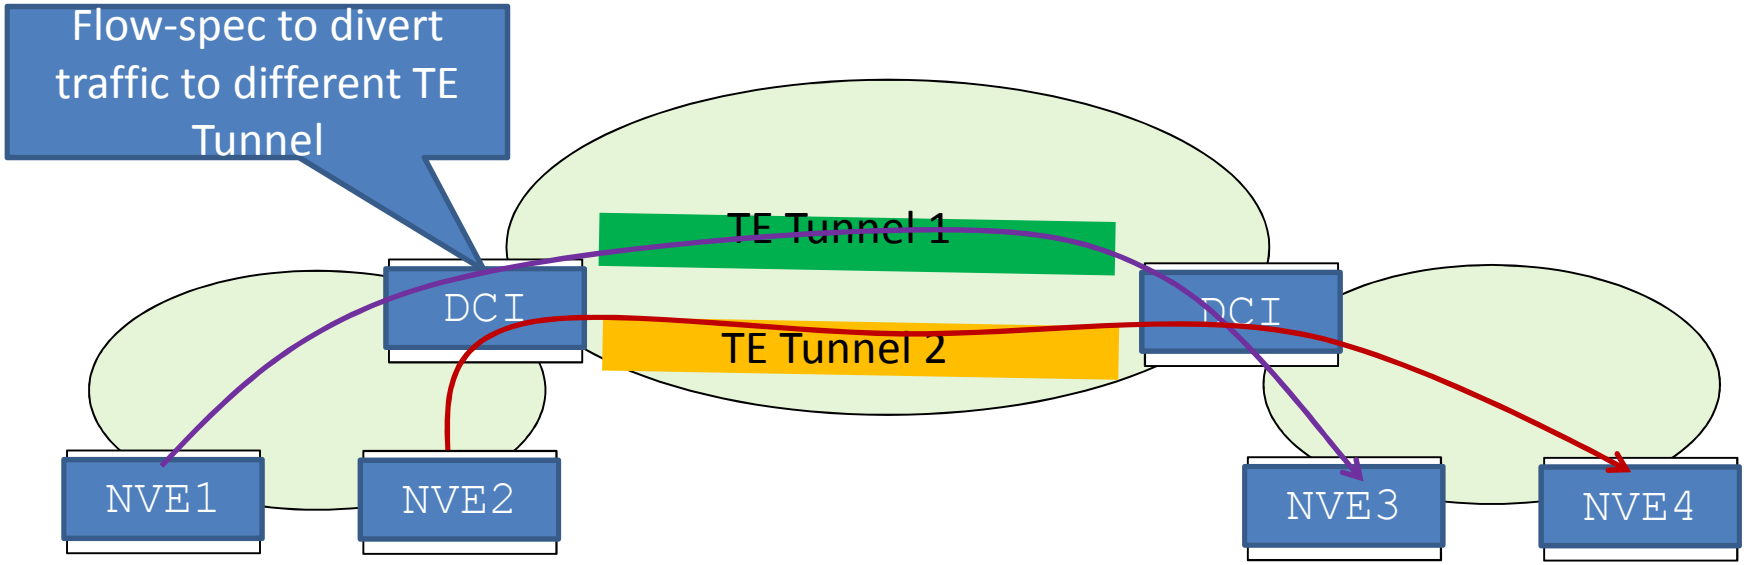


Divert traffic to different Tunnel relying on BGP flow-spec on DCI device.

# BGP Flow-spec for NVO3 Requirements Summary

① The match part should include inner L2/L3 header information and NVO3 header.

② The Traffic Filtering Actions supports redirect to TE tunnel or NVO3 tunnel.

# Flow-spec extension

Add new Component Types for NVO3 header:

```
Type TBD1 – VN ID
Type TBD2 – NVO3 Proto Type
```

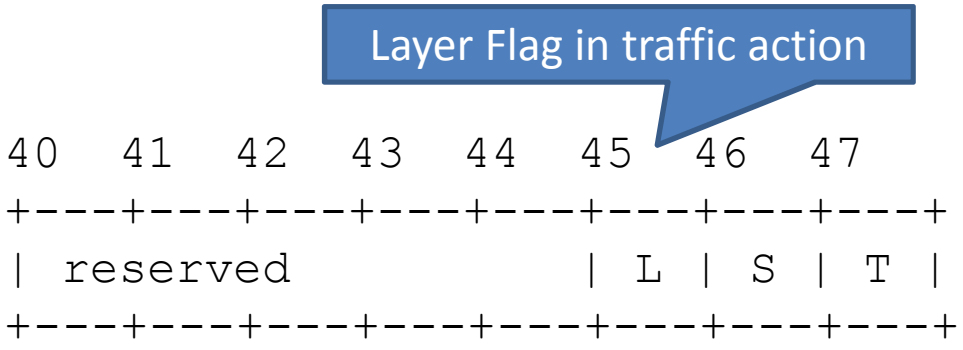


Layer Flag(Bit 45): When this bit is set, the corresponding filtering rules will be applied on the NVO3 inner layer. If not set, the corresponding filtering rules will be applied on the NVO3 outer layer.

# Next Step

- Seek some comments and feedbacks

# Flow-spec traffic action

| type | extended community | RFC or Draft |
|---|---|---|
| 0x8006 | traffic-rate | RFC5575 |
| 0x8007 | traffic-action | RFC5575 |
| 0x8008 | redirect | RFC5575 |
| 0x8009 | traffic-marking | RFC5575 |
| TBD | redirect to Tunnel | This draft |

- Besides ‘redirect to IP’ and ‘redirect to VRF’, ‘redirect to Tunnel’ is proposed.
- The set of tunnels can be specified in the BGP Path Attribute.

# BGP FlowSpec extensions for Routing Policy Distribution(RPD)

draft-li-idr-flowspec-rpd-00

Liang Ou, Yujia Luo
Robin Li, Vincent Zhuang, Eric Wu

IETF93, Prague

# Motivation

- Provider's motivation for traffic adjustment:

- Link congestion or overload;
- Packet delay, loss or corrupted;
- Prefer lower price;

# Genuine requirements

## Outbound traffic control

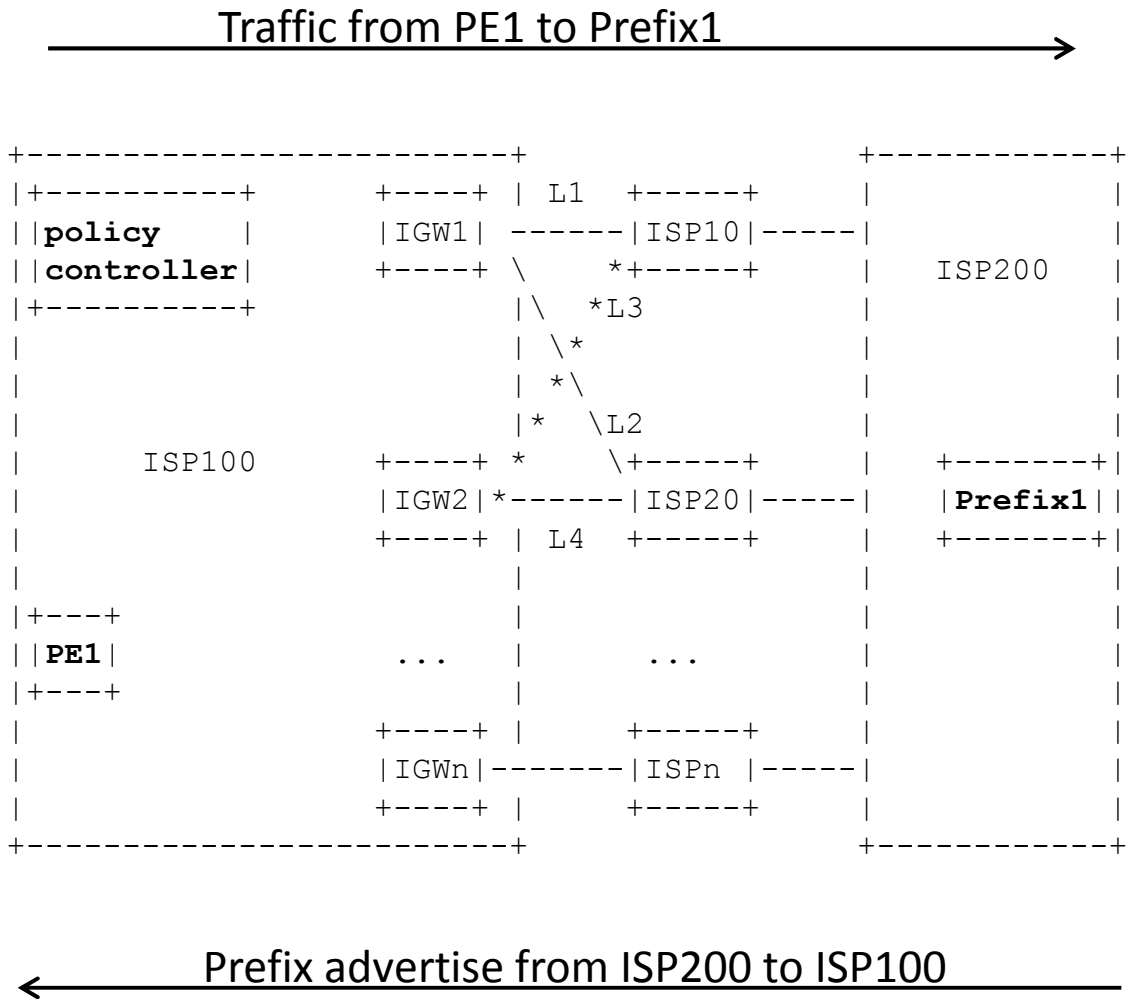


- EBGP peering:
  - IGW1---L1---ISP10
  - IGW1---L2---ISP20
  - IGW2---L3---ISP10
  - IGW2---L4---ISP20

- **Requirement:**
  - **Exit through L3**

# Genuine requirements

## Inbound traffic control

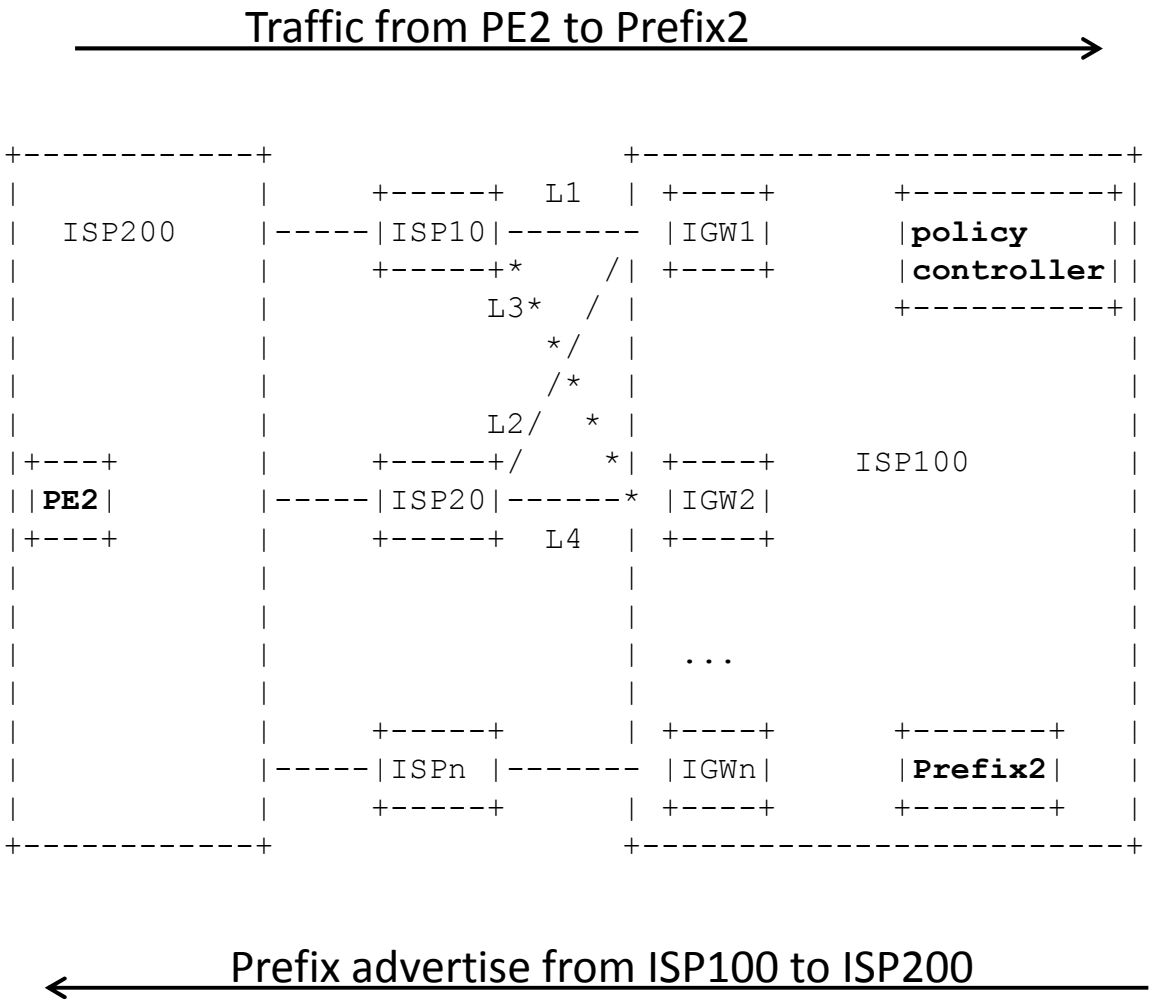


- EBGP peering:
  - IGW1---L1---ISP10
  - IGW1---L2---ISP20
  - IGW2---L3---ISP10
  - IGW2---L4---ISP20

- **Requirement:**
  - **Enter through L3**

# Solution requirements

- Minimal configuration burden, less manually;
- Dynamic provisioning instead of static one;
- Safe to undo or rollback without damage;
- Easy to maintain with reasonable complexity;
- High performance mechanism with small delay;
- Backward compatible with less software upgrade;

# Proposed solution

- Routing Policy Distribution(RPD)
  - Taking effect on control plane
  - Impact decision on remote site

- RPD protocol: BGP Flowspec
  - Filtering rule: destination for prefix1/prefix2
  - Action: R-bit introduced, more info carried in new attribute

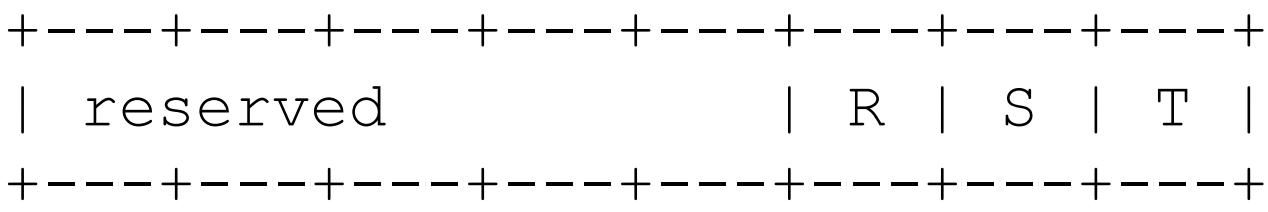

# Proposed solution

## ☐BGP Policy Attribute

- Attribute structure

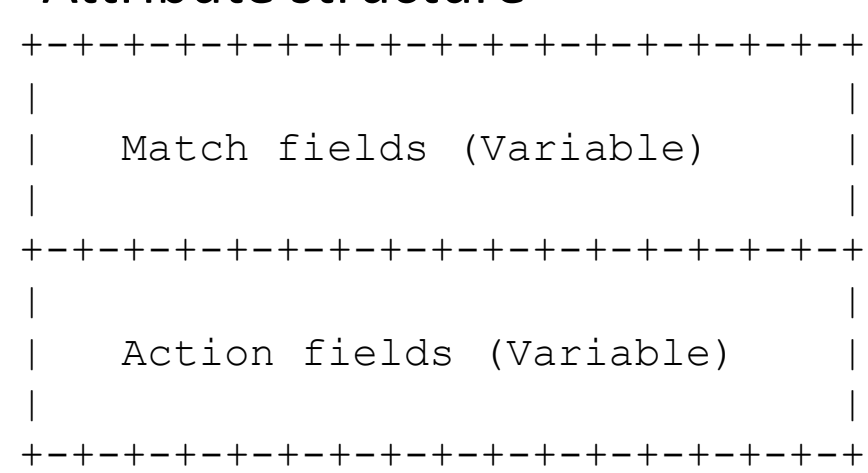


- Match field

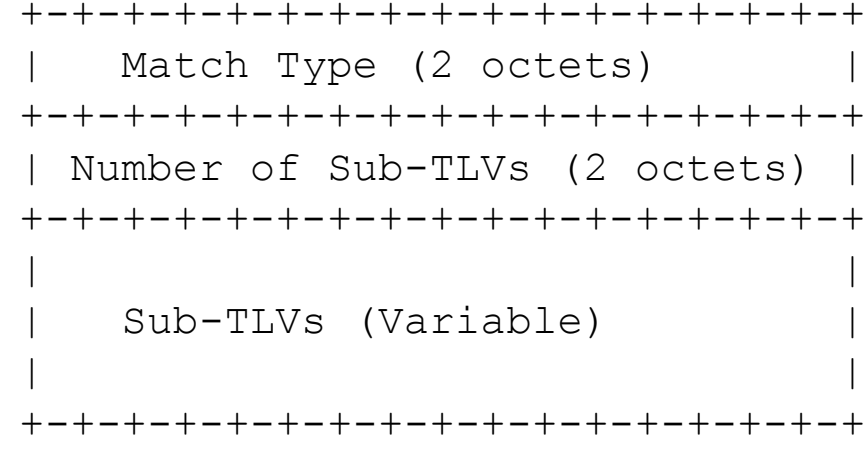


- Action field




- Action type 1: Route-Preference
- Action type 2: Route-Prepend-AS

☐Match type

- Value 0: Positive match
- Value 1: Negative match

☐Sub-TLVs

- Type 1: IPv4 Neighbor
- Type 2: IPv6 Neighbor
- Type 3: ASN list

# Proposed solution

- Outbound traffic control

- Match type: permit
- IPv4 neighbor sub-TLV:
  - ✓Local BGP Speaker IGW2
  - ✓Remote BGP Peer ISP10
- Action type: Route-Preference

- Inbound traffic control

- Match type: deny
- IPv4 neighbor sub-TLV:
  - ✓Local BGP Speaker IGW2
  - ✓Remote BGP Peer ISP10
- Action type: Route-Prepend-AS
- Action value: Prepend-AS five times

# Comparing to traditional routing policy

- Drawbacks using traditional routing policy:

- Manual configuration burden and mistakes.
- Complexity increased and difficulty to maintain.
- Static configuration versus dynamic requirements.

# Comparing to traditional BGP-Flowspec

## ☐Current BGP-Flowspec limitation

•Can't affect AS external device's decision.

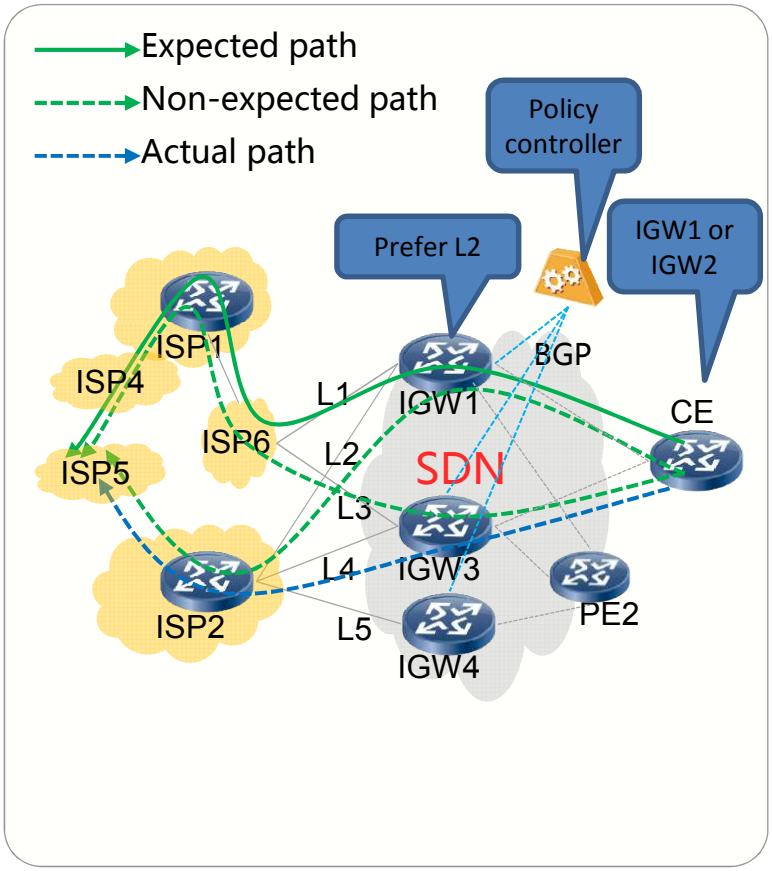


•On IGW1, L2 is preferred previously.

•Flowspec can redirect to L1

•CE prefers IGW3, which makes IGW1's redirection useless.

# Comparing to traditional BGP-Flowspec

## ☐Current BGP-FS limitation

- Changing decision in forwarding plane may introduce loop.

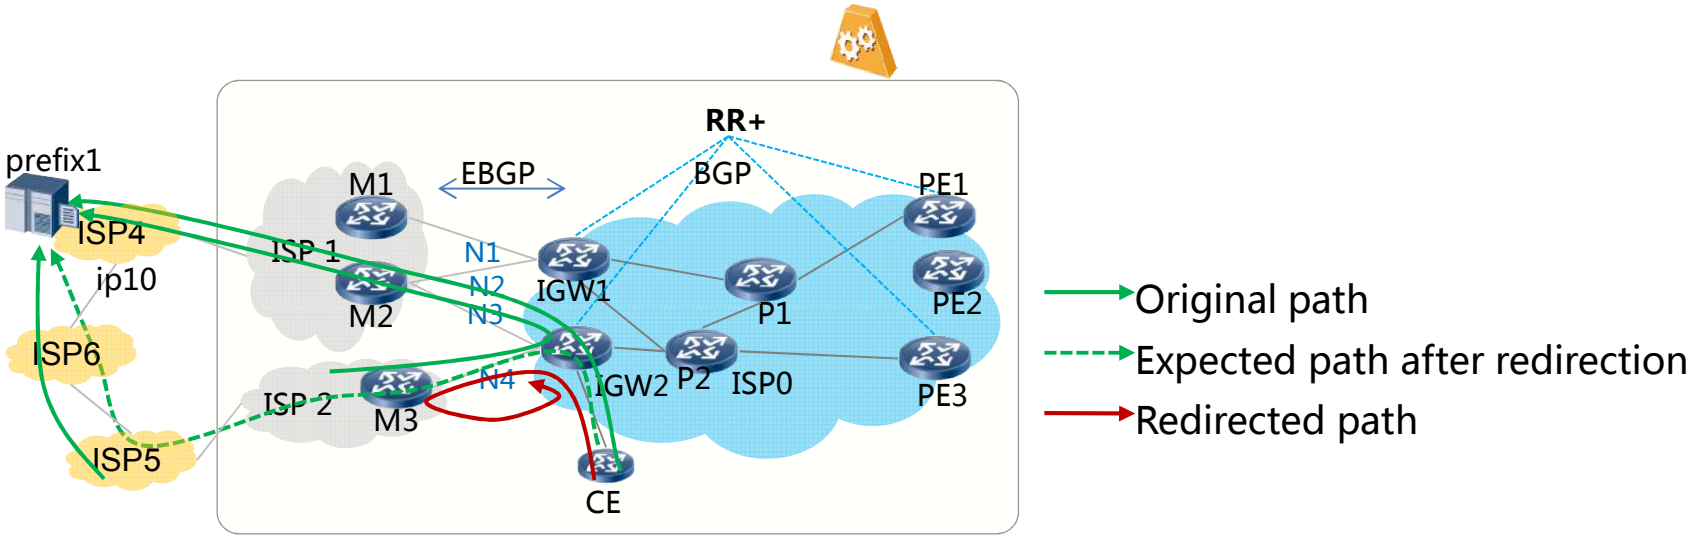


- IGW2 redirect from M2 to M3;
- M3 may still prefer IGW2's nexthop
- Loop can happen between M3 and IGW2.

# Next step

- ☐ Collect feedback and comments.
- ☐ Refine this draft according to comments.

# Label Information for BGP FlowSpec

## draft-liang-idr-bgp-flowspec-label-00

Jianjie You (youjianjie@huawei.com)
Qiandeng Liang (liangqiandeng@huawei.com)

IETF93 Prague

# Motivation

- Improve routing efficiency
  - FlowSpec rule mapped to a label
- Require traffic statistics per FlowSpec rule
  - Traffic statistics granularity should be improved based on FlowSpec rule rather than the destination prefix

# Scenario

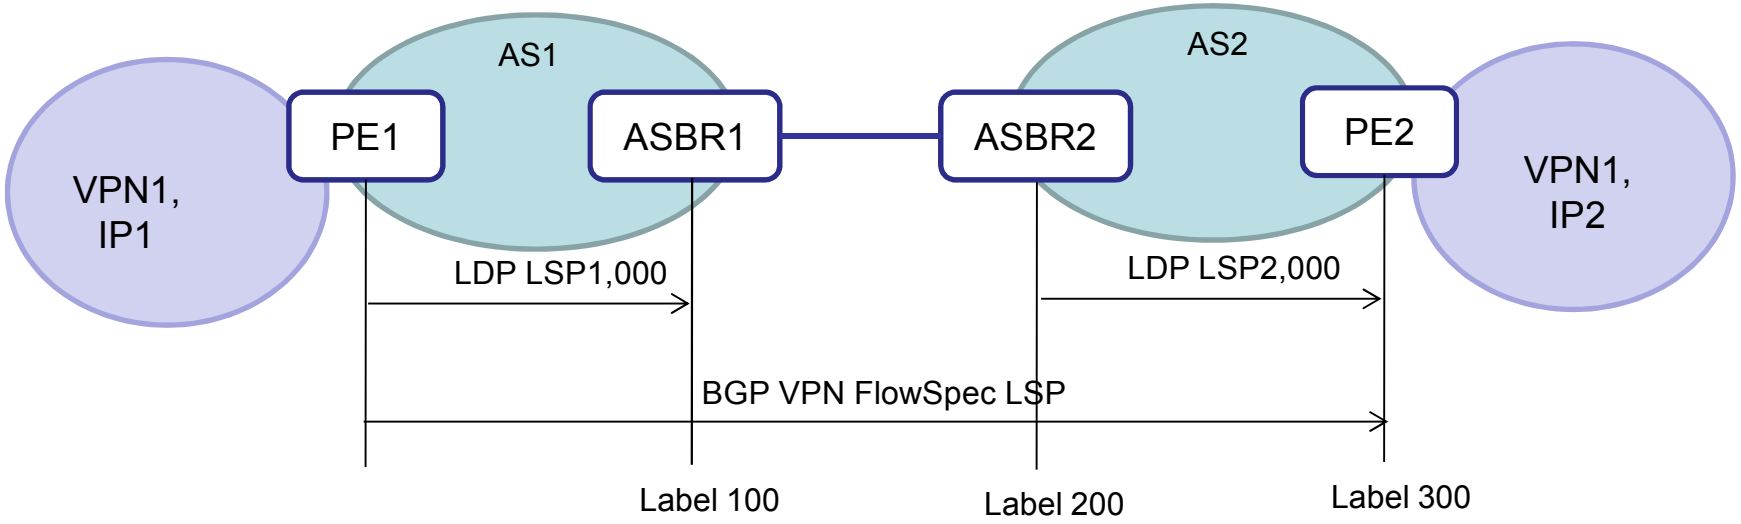


- FlowSpec Rule 1 (injected in PE2)
  - Filters: Destination IP prefix:IP2/32; Source IP prefix:IP1/32
  - Actions: traffic-marking: 1 (DSCP value)
- Forwarding Process on PE1 when receiving traffic from IP1 to IP2
  - PE1: Push 1,000 and 100
  - ASBR1: Pop 1,000, and then swap 100 to 200
  - ASBR2: swap 200 to 300, and then push 2,000
  - PE2: Pop all labels

# Solution 1: New Flow-spec label(s) route

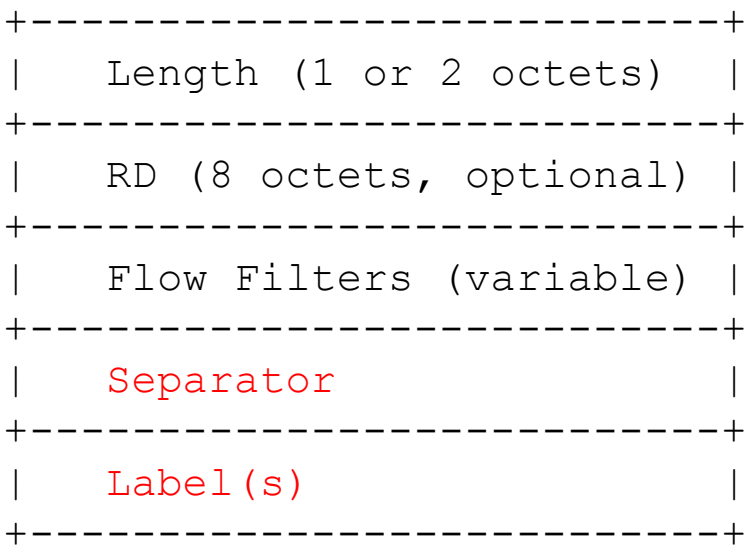


- **Separator**: indicates the separation between the flow filters field and lable(s) field. This value should not be used by flow filters.

- **Label(s):** This field carries zero or more labels (that corresponds to the stack of labels [RFC3032]).

# Solution 2: Using RFC3107 to carry FlowSpec label(s)

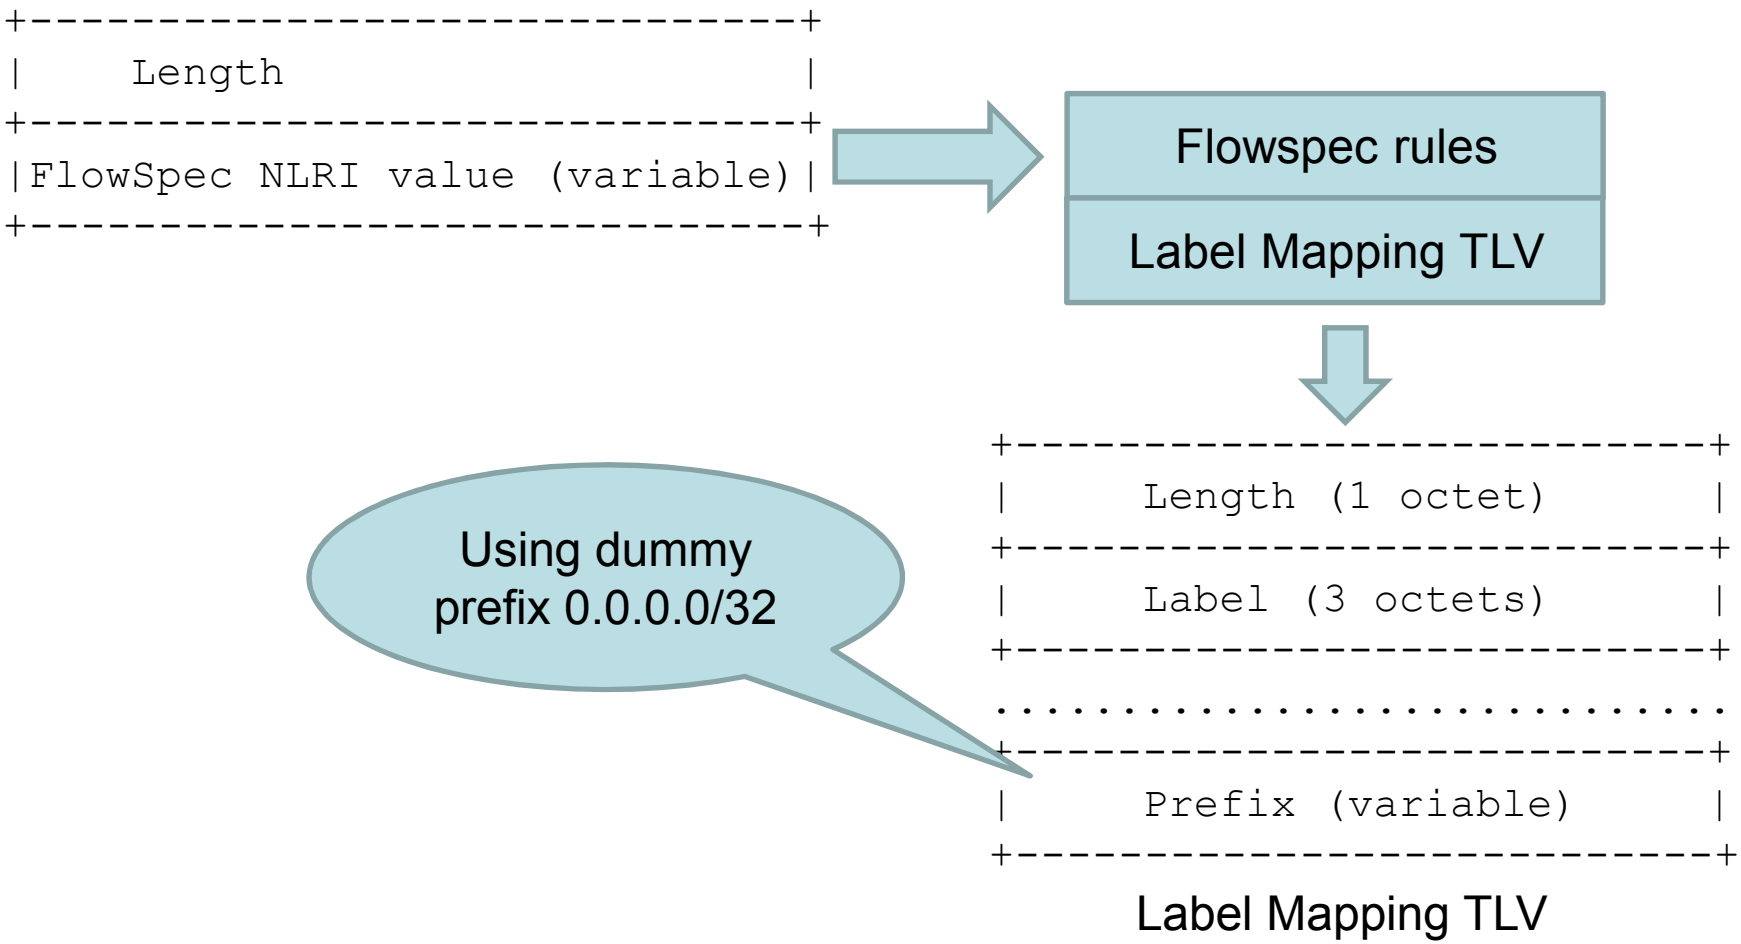

# Next Step

- Solicit comments and suggestions on the mailing list

# Thank You!

# Flow Specification Yang Model

https://tools.ietf.org/html/draft-wu-idr-flowspec-yang-cfg-01

Aseem Choudhary, Nan Wu, Shunwan Zhuang

*IETF 93, July 2015, Prague*

  

# Topics

- Background
- Flow Specification YANG Model Overview
- Flow Specification Model
- Flow Specification Configuration Data
- Flow Specification State Data
- Possible changes in next revision
- Questions

# Background

- This draft was earlier published as draft-wu-rtgwg-flowspec-yang-cfg-00
- Most of the work has been done for BGP flow specification
- Addressed a few configuration and state parameters issues and filled some gaps
- Republished as draft-wu-idr-flowspec-yang-cfg-00

# Flow Specification Yang Model Overview

- Defines Yang model for configuration and state data
- Models flow specification rules as defined in RFC 5575
- Models flow specification state data as route information and flow counters
- RFC 6087 Compliance
- Any vendor specific configuration or operational data is augmentable as vendor extensions

# Flow Specification Model

- Flowspec module has Flowspec as a top level container
- Flowspec container maintains Flowspec-cfg and Flowspec-state containers for configuration and state
- Flowspec-cfg maintains list of flow specification policies
- A Flowspec policy stores list of flow specification rules
- Flowspec-state maintains Flowspec-rib and Flowspec-statistics containers
- Flowspec-rib maintains list of flowspec routes
- Flowspec-statistics stores counters for flowspec rules

# Flow Specification Model

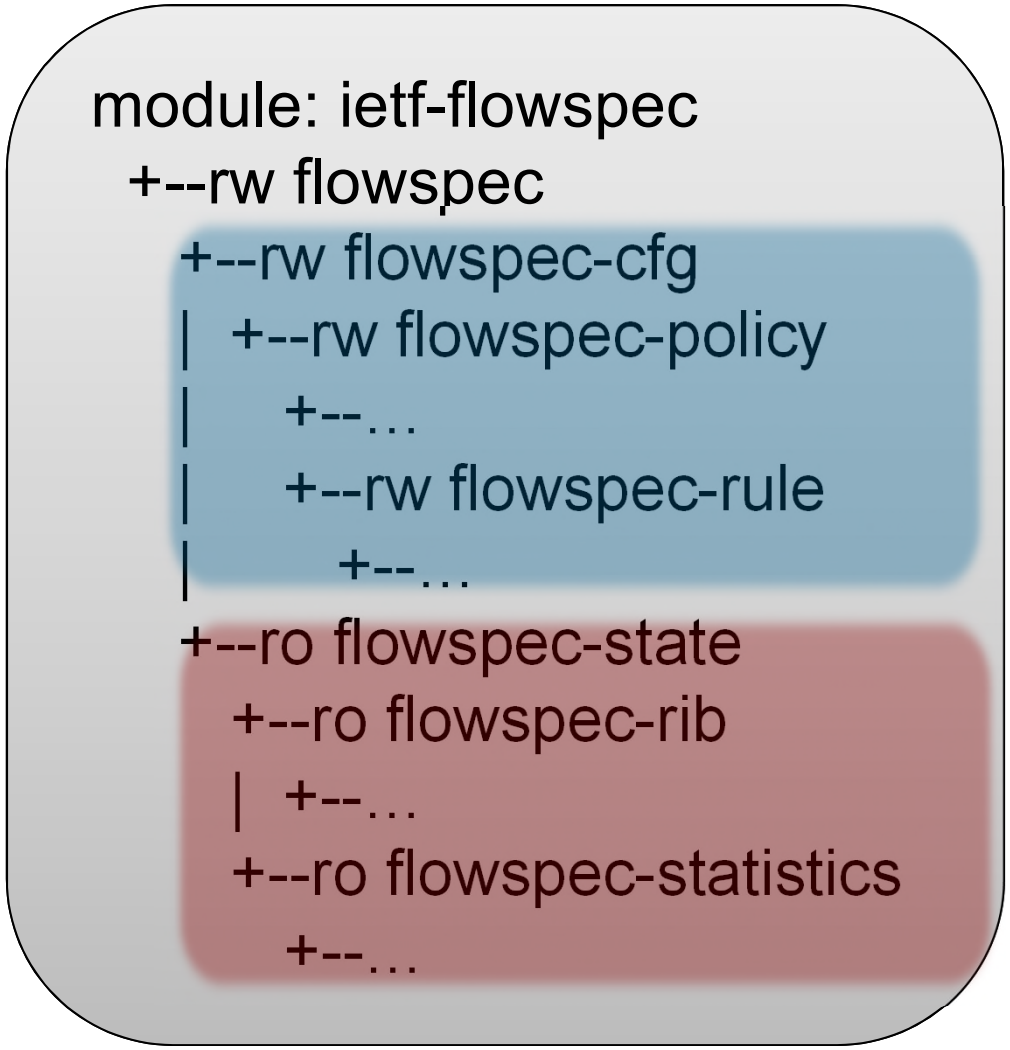


Configuration Container

State Container

# Flow Specification Configuration Data

- Flowspec policy is referred by name
- Flowspec policy is associated with VRF and address family
- Flowspec rule, is referred by name, contains list of components and actions
- Flowspec components referred by type
- Support destination IP prefix, source IP prefix, IP protocol, port, source port, destination port, ICMP type, ICMP code, packet length, TCP flags, DSCP, fragment
- Flowspec actions referred by action type
- Supports traffic-rate, redirect, traffic-marking

# Flow Specification Configuration Data

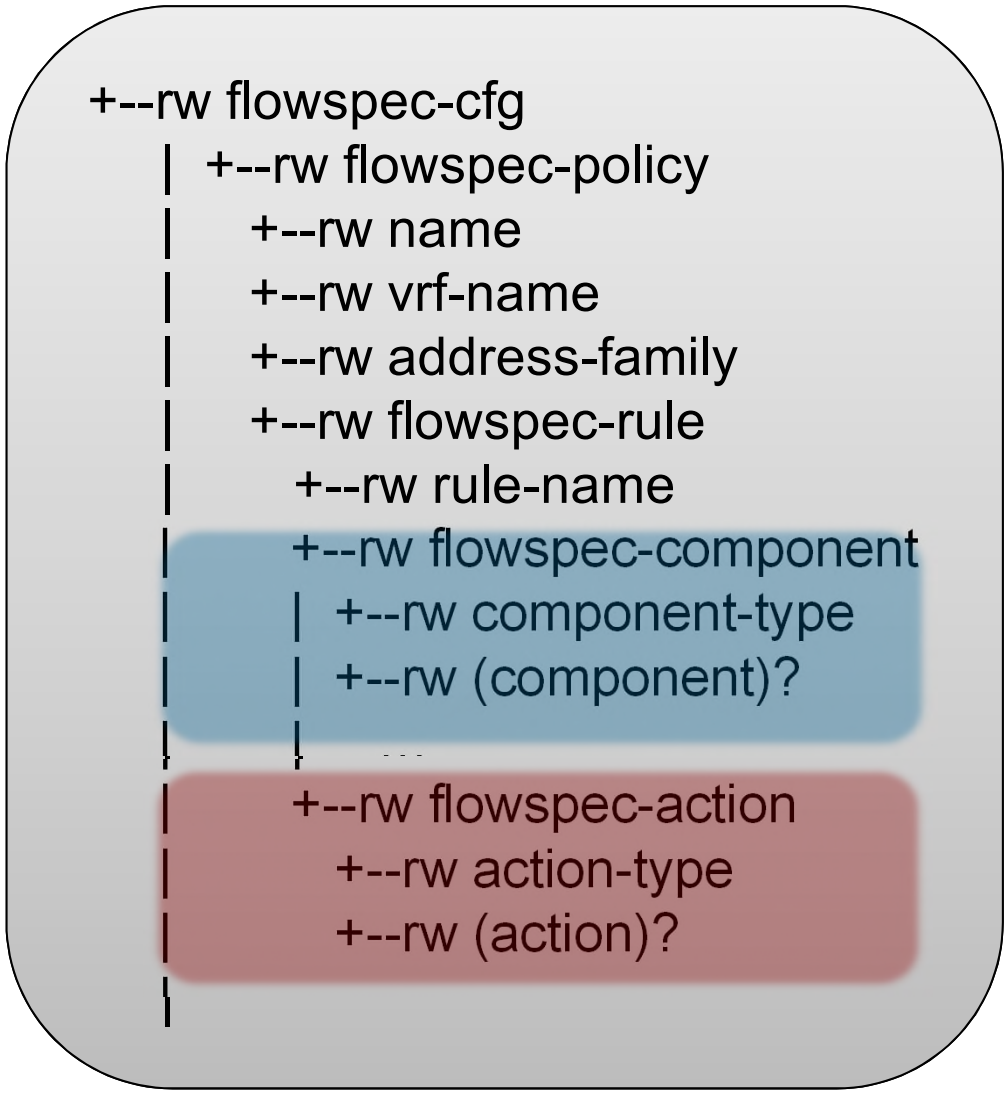


List of components

List of actions

List of rules

# Flow Specification State Data

- Flowspec-state maintains routes and flow counters
- Flowspec-rib container maintains generic as well as protocol specific data
- Flowspec rules are stored in protocol specific data
- Flowspec statistics container maintains flow counters per VRF and address-family
- Supported counters are classified packets, classified bytes, drop packets, drop bytes

# Flow Specification State Data

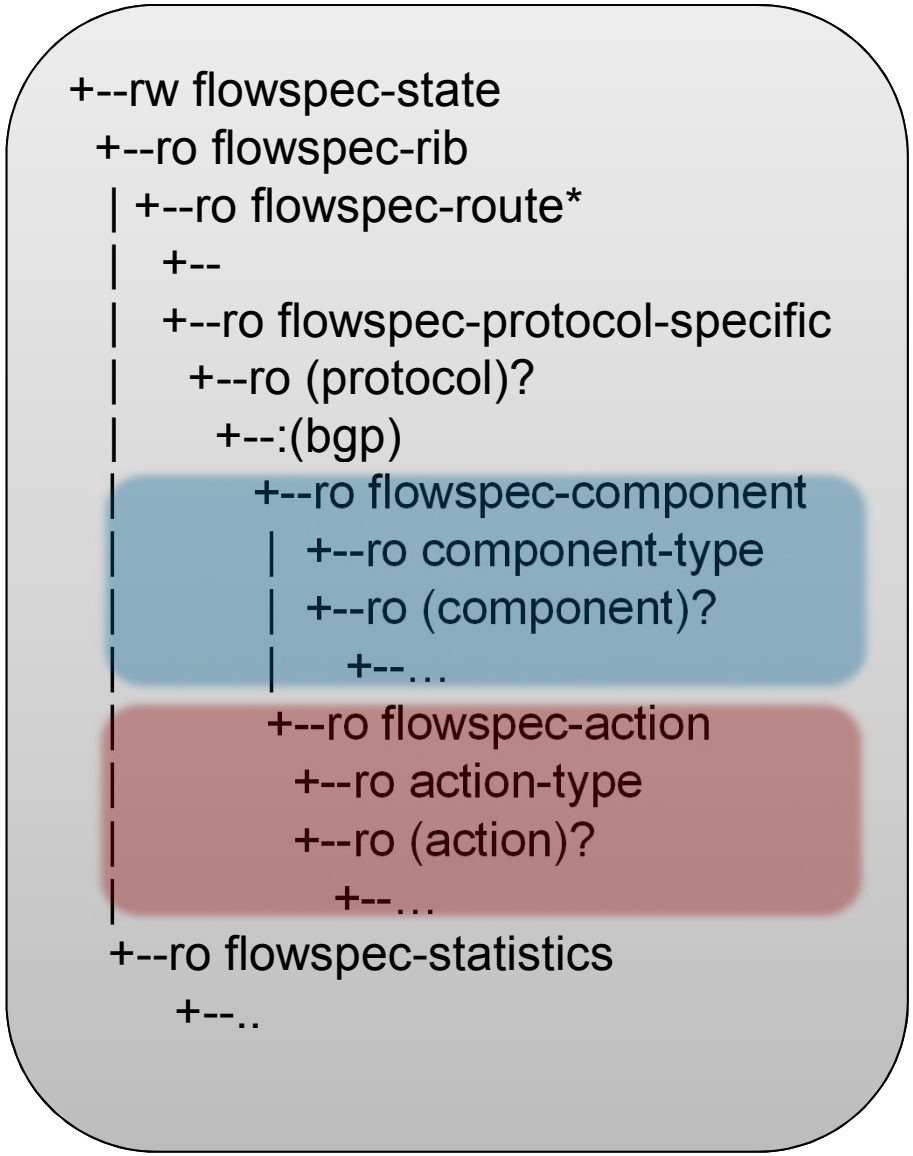


List of components

List of actions

# Flow Specification State Data

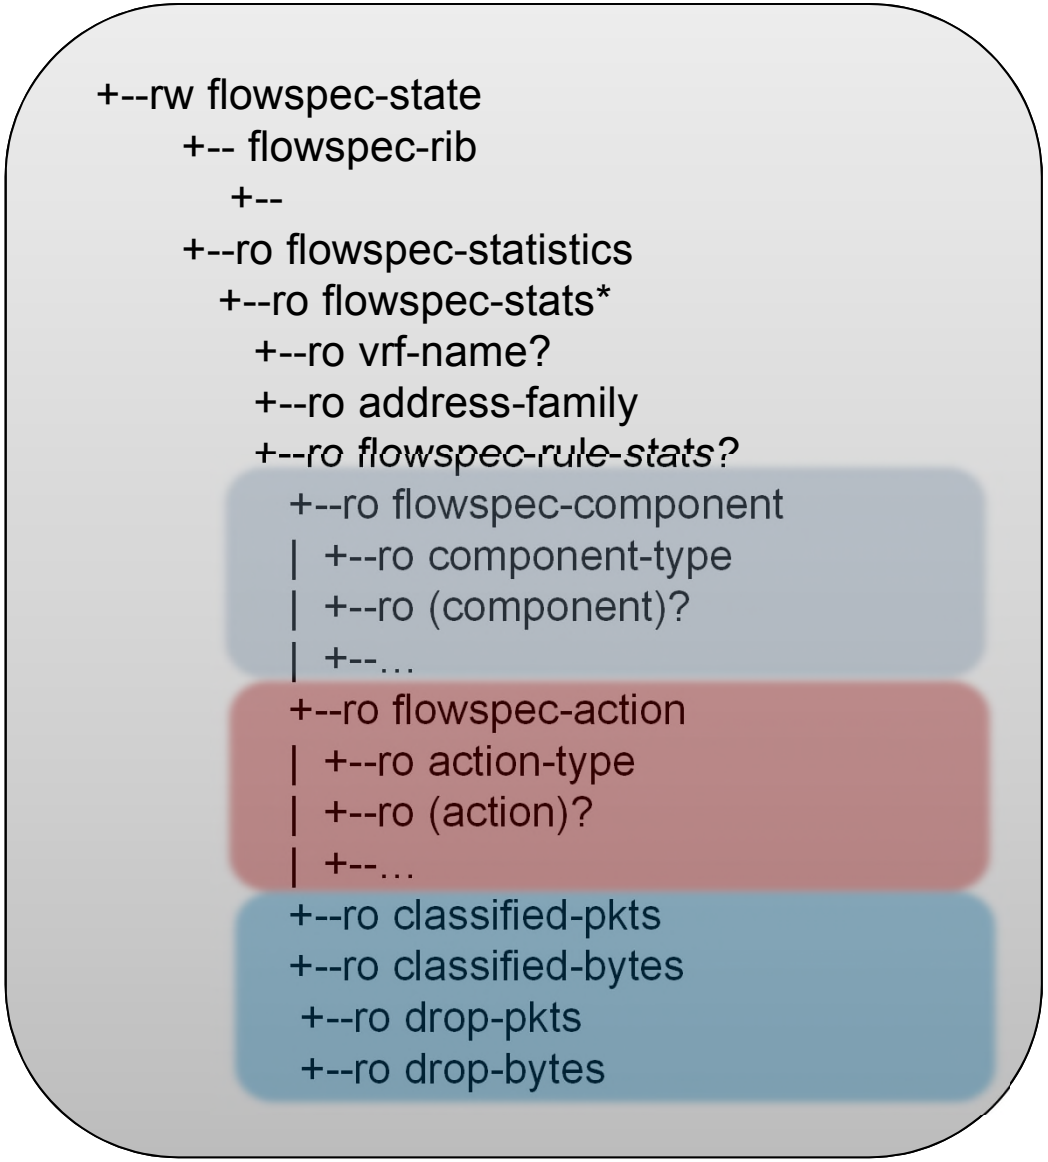


List of components

List of actions

Flow Counters

# Possible changes in the next revision

- Support of V6 Parameters
- Support other flow specification drafts in the model
- Changes due to review comments

Questions?

# BGP FlowSpec Outbound Route Filter

## draft-liang-idr-flowspec-orf-00

Qiandeng Liang (liangqiandeng@huawei.com)
Weiguo Hao (haoweiguo@huawei.com)
Jianjie You (youjianjie@huawei.com)

IETF93 Prague

# Motivation

- Network Security Policies
  - ASBR filters traffic from adjacent AS – reducing Traffic

- FlowSpec Specific Capability Negotiation
  - Select from all Types of ORFs to just the one's selected

# Scenario

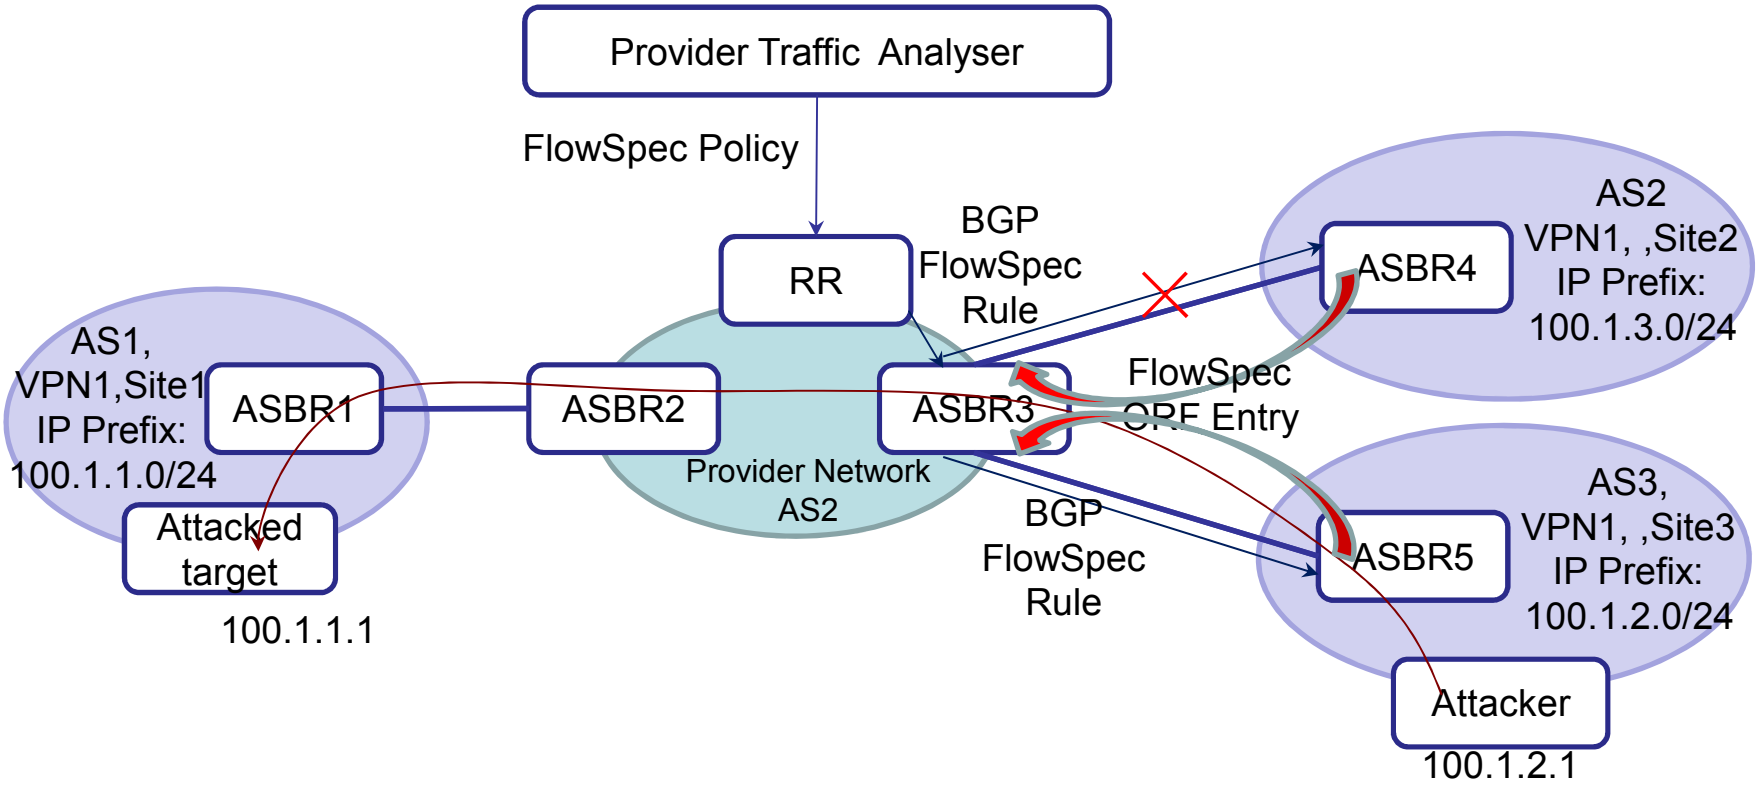


VPN1,Site1 only has an internal subnet(IP prefix:100.1.1.0/24).
VPN1,Site2 only has an internal subnet(IP prefix:100.1.3.0/24).
VPN1,Site3 only has an internal subnet(IP prefix:100.1.2.0/24).

Attacker with IP address as 100.1.2.1 in the network of VPN1,Site3,
who is attacking the host 100.1.1.1 in the network of VPN1,Site1.
And the traffic analyzer has detected this attacking traffic .

.

# Protocol Extensions(1)

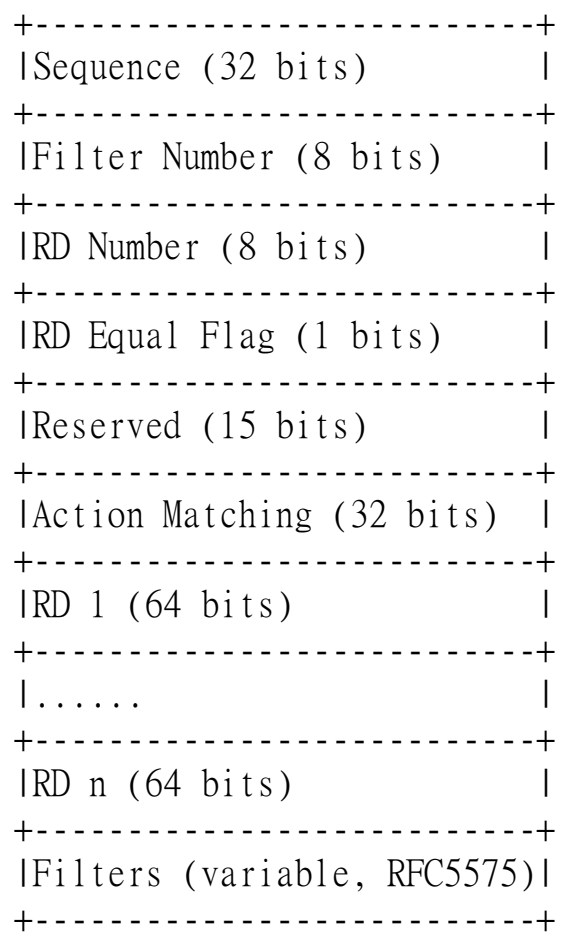


FlowSpec-ORF
"Type specific part" Encoding

- **RD Number:** the number of the RD items.
SAFI=133, RD Number=0;
- SAFI=134, RD > 1

- **RD Equal Flag:** matching mode for RD.

- **RD:** An 8-byte Route Distinguisher (RD), present when SAFI=134.

- **Action Matching:** each bit corresponds to a particular FlowSpec action [RFC5575].
- If set, match the action;
- If unset, not match the action.

# Protocol Extensions(2)

The filters in FlowSpec-ORF are aligned with the filters defined in [RFC5575], etc. except the following four types:

| Filter Type | RFC/Draft |
|---|---|
| Type 1: Destination Prefix IPv4 or IPv6 | RFC5575, draft-ietf-idr-flow-spce-v6 |
| Type 2: Source Prefix IPv4 or IPv6 | RFC5575, draft-ietf-idr-flow-spce-v6 |
| Type 14: Destination MAC Prefix | draft-hao-idr-flowspec-evpn |
| Type 15: Source MAC Prefix | draft-hao-idr-flowspec-evpn |

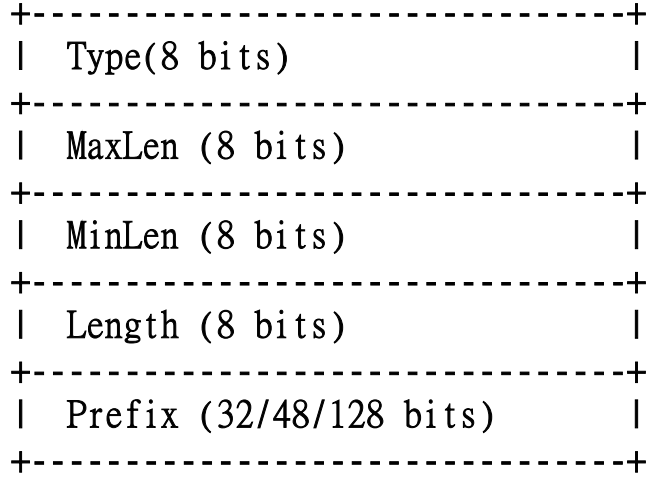


These four types are encoded as shown in the left figure . MaxLen, MinLen, Length, Prefix are the same as defined in [RFC5292].

# Next Step

- Describe the regulation how flowspec-orf entries filter the flowspec rule more clearly.

- Solicit more discussion on the mailing list

# Thank You!

# Upstream Assigned Label Collision Solution

draft-zhang-idr-upstream-label-collision-solution-00

Sandy Zhang (ZTE Corporation)

Ying Cheng (China Unicom)

# Problem Statement

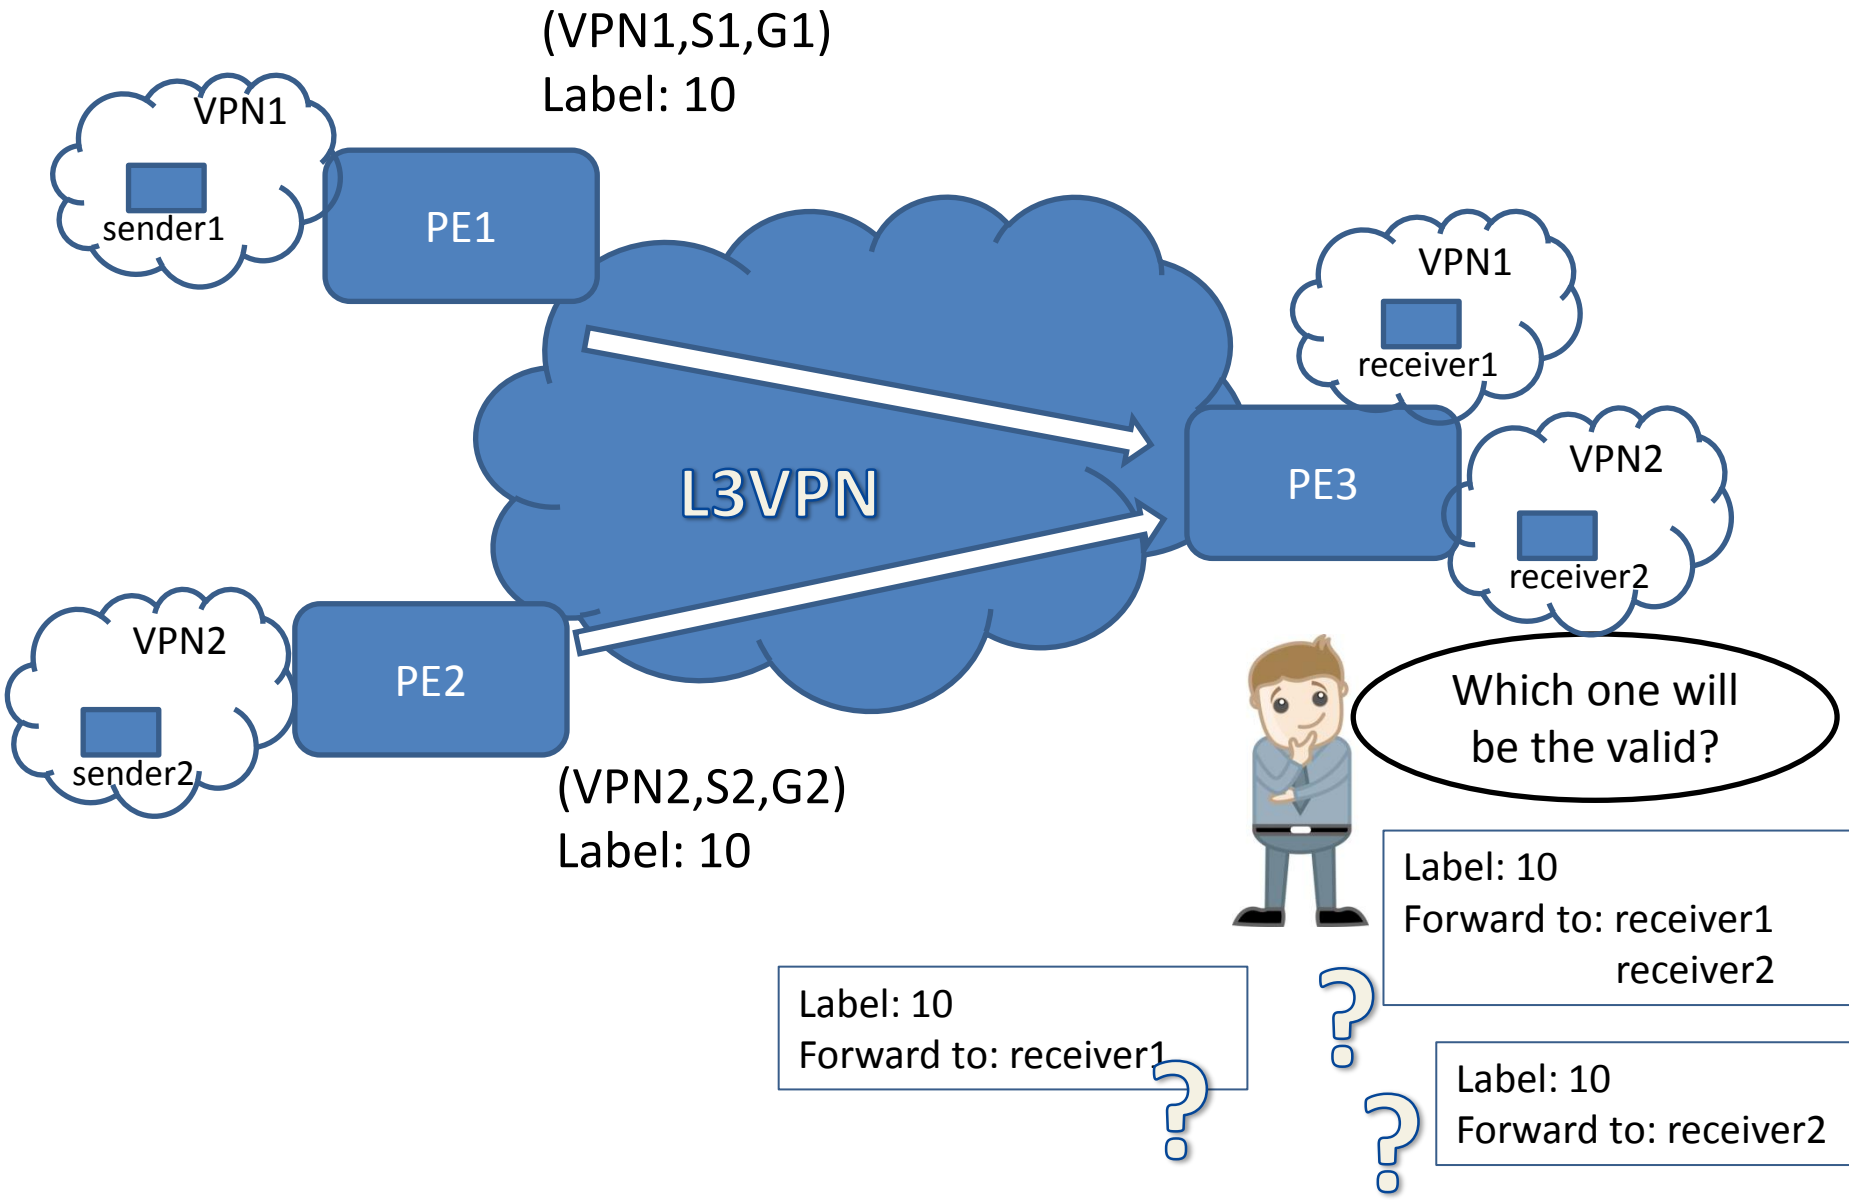

# Problem Statement

## Why do the collision happen?

- In present network, the network administrator must allocate the label space on every PE in advance.
- There are tens of PEs in one domain.
- The network is developing. The MVPN and DC will consume a great number of upstream-assigned labels.

-------The numbers of MVPN S-PMSIs will become larger.

-------The BUM of DC may consume more upstream-assigned labels.

- If the label space that is reserved on every PE is large, then many labels may be waste.
- If the label space that is reserved on every PE is small, then some PEs will use up all its labels.
- The network administrator adjust the label space on every upstream PEs. The adjustment is not foreseeable and consume large manpower.

# Solution

The algorithm:

- Each upstream PE advertise routes with the timestamp attribute.
- When downstream PEs finds that there are two upstream PEs advertise routes with a same upstream-assigned label:

------The downstream PEs choose the route with the earlier timestamp to be valid.

- All the PEs include upstream and downstream PEs will receive the routes.

------The upstream PEs that advertise route with later timestamp will adjust the label.

- Tie-Break

------If two upstream PEs advertise route with same timestamp, we make the tie-break on the IP-address of upstream PEs.

# Advantage of solution

- Why we use timestamp to describe the route?

------Someone may think that we may use the local-preference/MED/cost and so on.

- The timestamp will mark the origination time of the routes.
- The routes which have been originated early should have more priority.
- The method will improve the network stability.

- Comments welcome

Thanks!